

# **Sudharshan S. Vazhkudai**

## *Curriculum Vitae*

### Table of Contents

1.	Current Positions	2
2.	Profile	2
3.	Educational Experience	2
4.	Work Experience	2
5.	Projects	3
6.	Technical Leadership Experience	4
7.	Publications and Presentations	5
8.	Program Development Activities and Grants	9
9.	Awards	9
10.	Professional Activities	9
11.	Personnel Management and Advisees	9

---

## Sudharshan S. Vazhkudai

[vazhkudaiss@ornl.gov](mailto:vazhkudaiss@ornl.gov)  
865-576-5547  
<http://www.csm.ornl.gov/~vazhkuda>

Oak Ridge National Laboratory  
1 Bethel Valley Road  
Oak Ridge, TN 37831

---

### 1. CURRENT POSITIONS

**Group Leader**, Technology Integration, National Center for Computational Sciences, Oak Ridge National Laboratory (ORNL)

**Joint Faculty Associate Professor**, The University of Tennessee, Knoxville (UTK)

### 2. PROFILE

- More than 10 years of experience in the US Department Of Energy's (DOE) national lab system (nine at ORNL and two at Argonne National Lab, ANL), working in areas such as: HPC, distributed storage, large-scale data management/analytics, parallel file systems, NVRAM, multicores, runtime, and grids.
- A doctorate in computer science with a specialization in wide-area data management.
- Leading a multi-institutional research project, funded by the National Science Foundation (NSF) High-End Computing File Systems and I/O (HECFSIO) program, to investigate the utility of NVRAM as a staging storage and memory extension for supercomputers.
- From 2003 to 2010, led the design and development of distributed data management, distributed job orchestration and data archival workflow for the Spallation Neutron Source (SNS), a billion dollar infrastructure that produces O(100 TB) of scientific data.
- Technical lead for a National Institutes of Health (NIH) project to create a distributed data repository for manipulating and accessing oncology data, emanating from the National Cancer Institute sites.
- Led several projects on parallel file systems, deep-memory tiers, distributed storage, and data movement.
- Research contributions and tools for the petascale Jaguar machine (#6 in Top500).
- Obtained over four million dollars in research funding as principal investigator (PI) and Co-PI on grants from NSF, NIH, ORNL's Lab Directed R&D program (LDRD) and SEED programs.
- 50+ papers in highly selective conferences/journals, with around 1100 citations and an H-index of 13.

### 3. EDUCATIONAL EXPERIENCE

**Doctor of Philosophy** in Computer Science, May 2003

*Institution:* University of Mississippi/Argonne National Laboratory

*Dissertation:* Bulk Data Transfer Forecasts and the Implications to Grid Scheduling

*Highlight:* Wallace Givens fellowship from ANL; Thesis work on the Globus grid toolkit.

**Master of Science** in Computer Science, December 1998

*Institution:* University of Mississippi

*Thesis:* Performance Oriented Distributed OS-Evolutionary Steps towards a Distributed Linux

*Highlight:* Led a team of graduate students on the development of a Distributed Linux.

**Bachelor of Engineering** in Computer Science, June 1996

*Institution:* Karnatak University, India

### 4. WORK EXPERIENCE

#### OAK RIDGE NATIONAL LABORATORY

**Oak Ridge, TN**

11/2012 – present *Group Leader, Technology Integration, National Center for Computational Sciences (NCCS)*

2003 – 10/2012 *Research Staff Member, Computer Science and Mathematics Division (CSMD)*

#### THE UNIVERSITY OF TENNESSEE

**Knoxville, TN**

2010 – present *Joint Faculty, Joint Institute of Computational Sciences (JICS)/CEE*

2005 – 2006 *Adjunct Assistant Professor, Computer Science*

#### ARGONNE NATIONAL LABORATORY

**Argonne, IL**

2000 – 2002 *Givens Fellow/Doctoral Fellowship, Math and Computer Science Division*

#### THE UNIVERSITY OF MISSISSIPPI

**Oxford, MS**

1997 – 2000 *Instructor/Research Assistant, Department of Computer Science*

#### ESRI

**Redlands, CA**

Summer 1998     *Software Engineering Intern*  
**WIPRO INFOTECH**  
1996 – 1997     *Design Engineer, R&D*

**Bangalore, India**

## 5. PROJECTS

**Non-Volatile Memory (NVRAM):** The following projects highlight my experience in NVRAM technologies from three perspectives: fault tolerance, memory extension, and active computation. (i) **Distributed NVRAM-based Checkpoint Storage:** (2010-2013) Currently serving as a PI on a multi-institutional *NSF HECFSIO* project to build an aggregate, intermediate storage for supercomputers, using SSDs/DRAM, in order to alleviate the I/O bottleneck in checkpointing [18, 31, 56, 57] and improve fault tolerance. The checkpoint storage can be built by aggregating compute SSDs, DRAM or node-local disks (when available) and can be presented as an intermediate storage between massively parallel applications and the parallel file system (PFS) in supercomputing or desktop grid environments, respectively. Built novel incremental checkpointing support within the storage using checksum comparisons of chunks. *Team:* Two graduate students, a postdoc and three faculty collaborators. The project has contributed to one doctoral and one master's thesis. (ii) **NVMalloc for Out-of-Core Data Analytics:** (2011-) Leading an NSF effort to use an aggregate NVRAM store as a secondary memory partition for out-of-core data analysis [15]. The NVMalloc library helps applications conduct data analytics using NVRAM allocations instead of DRAM. Built: (a) a runtime library for applications to explicitly allocate and access variables on a distributed NVRAM store; (b) techniques to map a byte-addressable interface to a block store and (c) seamless checkpointing of DRAM and NVRAM allocated variables. Currently investigating statistical techniques for smart program variable (data) placement across deep-memory tiers. *Team:* A graduate student, a postdoc and two faculty collaborators. (iii) **Active Flash for “true” Out-of-Core Data Analytics:** (2011-) Leading an effort to expedite data analysis pipelines by migrating them to the location of the data, the flash device itself [12, 13, 16, 17], thereby freeing up both the compute core and the associated DRAM. Our work is an exploration into the following: performance/energy modeling of the various data analysis workflow scenarios seen in HPC centers, the ability of embedded controllers (ARM processors) to perform analysis alongside an HPC application run, a simulation study of Active Flash scheduling policies, the suitability of an analytics kernel for out-of-core execution using different fidelity modes, and prototype implementation using the OpenSSD platform. *Team:* A graduate student, intern and two faculty collaborators. This work is contributing towards a Ph.D. thesis of a graduate student.

**Parallel File systems (PFS):** The following projects highlight my work in PFS from the following perspectives: scratch PFS as a cache, recovery, contention, and end-to-end data integrity. (i) (2007-2009) Led an effort (Co-PI on an *NSF HECFSIO* proposal) to view the *HPC center's scratch storage as a cache* [26]. The motivating factor here is the observation that while a scratch PFS deals with a job input/output workload, similar to a cache workload, it is treated as a typical file system. Developed a coordinated scheduling mechanism that coincides data staging, computation and data offloading in an attempt to view the PFS as a cache (tool deployed on ORNL's 3.3 petaflop Jaguar machine and at Boeing's HPC center; it has been used by several applications.) *Team:* A graduate student, faculty collaborator and two professional staff. (ii) Led the development of novel *semantics for recovering transient job input data* on scratch PFS. These include *temporal replication* and *data failover* (to remote source copy) strategies for job input data for Lustre [6, 25, 30, 34]. *Team:* Two graduate students and two faculty collaborators. These two projects together have contributed towards three students' doctoral thesis work. (iii) (2012-) Co-PI on a project to *analyze the contention on shared storage systems* (e.g., the scratch PFS on supercomputers), develop mechanisms to facilitate HPC center-wide I/O operation coordination, and I/O-aware job scheduling. The proposed research will enable applications to coordinate their I/O operations at runtime to reduce inter application I/O interference, based on real-time load information from the underlying shared storage systems. *Team:* A graduate student and two staff members. (iv) (2011-) Contributing to an effort to use *GPUs to achieve RAID-6* for Lustre [14]. Our system utilizes low-cost, strategically placed GPUs—both on the client and server side—in a cluster to accelerate parity computation, enable fault-tolerance on a per file basis, provide end-to-end data integrity checks and parallelization of RAID array reconstruction. *Team:* A graduate student and two faculty collaborators.

**FreeLoader Cache using Distributed Storage Scavenging:** (2003-2007) Led the design and development of an aggregated, distributed storage infrastructure to be used as a client-side cache by data-intensive

applications [3, 8, 35, 38, 39, 41, 42, 43, 58, 59]. The basic idea is the aggregation of space and I/O bandwidth contributions from commodity desktops within a domain to provide a mountable (through FUSE), highly-available, shared cache/scratch space for large, immutable scientific data sets that can be accessed in parallel. Built novel techniques such as: asymmetric striping, prefix caching with suffix patching from remote sources, cost-of-recovery of a dataset to determine its redundancy scheme, etc. The suffix patching technique was even applied in the context of Lustre as a data failover strategy highlighted earlier. *Team*: Several graduate students, interns, post-masters staff and two faculty collaborators. The work has resulted in three master's theses.

**Functional Partitioning Runtime for Many-core Systems:** (2010-) Leading an effort to functionally partition many-core systems ( $O(100)$  cores/node) to optimize effective application performance [20]. Developing a novel runtime environment to dedicate cores to certain services needed to support a large-scale application run. Using this new paradigm, both critical exascale and end-to-end application assist tasks can be scheduled in-situ, alongside the application's simulation component for better end-to-end performance. For example, one can imagine spatially multiplexing data analytics/reduction routines, resilience services, deep-memory data placement, monitoring, etc., alongside the application's simulation, on the same node, in an attempt to improve the application execution. *Team*: Two graduate students and two faculty collaborators. This work is part of a graduate student's master's thesis.

**Wide-area Data Movement:** (i) (2007-2009) PI on a project to expedite the *end-user data delivery* between HPC centers and end-user locations [1, 2, 4, 19, 23, 28, 32, 33]. The premise here is that while HPC centers themselves are well connected, end-users are stymied by the last-mile problem. Developed decentralized data-offloading and just-in-time staging schemes to move job output and input data by reconciling center purge policies, user delivery and job startup deadlines. Combined point-to-point transfer tools (e.g., GridFTP) and decentralized schemes (e.g., BitTorrent) along with NWS to enable timely data delivery. Using such a scheme, users can exploit advanced networks (100Gb/s) and commodity networks in a seamless fashion to deliver data in time. Work resulted in a graduate student's Ph.D. thesis. *Team*: A graduate student and a faculty collaborator. (ii) (2000-2003) As part of my doctoral work, I developed middleware for the orchestration of bulk data transfers in the *Globus Data Grid*. The middleware comprised of a scalable storage broker and selection heuristics for locating widely replicated data, statistical models and tools to predict the performance of wide-area data transfer times, and techniques for co-allocating transfers [9, 10, 45, 46, 47, 48, 49, 52, 60, 61].

**Scientific Data Management for Application Domains:** (i) Between 2003 and 2010, I was the lead architect for a scientific gateway, a distributed computation solution and data archival management [7, 21, 22, 24, 27, 29, 36, 40] for  $O(100)$  TB from the **Spallation Neutron Source (SNS)**. Developed: (a) a science gateway solution for remote access to SNS data, computation, viz and proposal management; (b) data management strategies for handling hundreds of TBs of data: delivery, orchestration for user jobs, remote workspaces, scientific data (NeXus) browsing, automatic archival/retrieval/reconstruction to/from HPSS, community account (*jimmyneutron*) based data access and quality assurance; (c) application management strategies for seamless remote executions on a variety of resource managers: SLURM, PBS (OIC), Globus (TeraGrid.) *Thousands of portal visits, several TBs of data delivered and thousands of jobs are being run through this infrastructure.* (ii) **Oncology Data Coordination Center:** (2012-) Technical lead and Co-PI on a newly funded *NIH effort* to build a data repository at University of Tennessee so that National Cancer Institute (NCI) funded centers can aggregate, annotate, search, index, and download data.

**Cloud Storage:** (2010-2011) Led the development of a FUSE file system atop Azure cloud storage. We have developed mechanisms to use the cloud storage as intermediate nodes for data delivery from HPC centers to end-users [19]. *Team*: An intern and a staff member.

**Distributed OS for Linux:** (1998-2000) Built a high-speed communication protocol for a Linux cluster by short-circuiting the protocol stack; a networked file system, global IPC mechanism, group communication and remote process execution environment using the communication scheme [11, 53, 55, 62]. Patches for the 2.0 kernel. *Team*: Led several graduate students.

## 6. TECHNICAL LEADERSHIP EXPERIENCE

I have led several multi-institutional projects comprising of several graduate students, post-masters, postdoc, and university faculty researchers as illustrated above. The FreeLoader scientific data cache and the NSF project on distributed SSD-based checkpoint storage system and NVMMalloc are by far the largest projects I have led, comprising of 5-6 team members. In all of the aforementioned projects, I have been responsible for setting broad research directions, assigning weekly tasks, mentoring, monitoring progress, and publishing papers. In addition, in many cases, I am also involved in problem solving, coding, and debugging as needed. I have trained and mentored several students, postdoctoral and professional staff. I have served on several masters and doctoral committees of graduate students from North Carolina State University, Virginia Tech, Pennsylvania State University, and Northeastern University. For my educational outreach, I was awarded the “*Outstanding Mentor Award*” by the Oak Ridge Institute for Science and Education in 2008. A complete list of all of my advisees can be found at the end. In addition, I am currently the technical lead for an effort to create a large-scale data repository for cancer data. This involves researchers from several divisions at ORNL.

## **7. PUBLICATIONS and PRESENTATIONS**

### **Journals**

1. H. Monti, A.R. Butt, S.S. Vazhkudai, “On Timely Staging of HPC Job Input Data,” *to appear in the IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2012.
2. H. Monti, A.R. Butt, S.S. Vazhkudai, “Timely Result-Data Offloading for Improved HPC Center Scratch Provisioning and Serviceability,” *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Vol. 22, No. 8, pp. 1307-1322, August 2011.
3. X. Ma, S.S. Vazhkudai, Z. Zhang, “Improving Data Availability for Better Access Performance: A Study on Caching Scientific Data on Distributed Desktop Workstations,” *Journal of Grid Computing - Special Issue on Volunteer Computing and Desktop Grids*, Vol. 7, No. 4, pp. 419-438, December 2009.
4. S.A. Kiswany, M. Ripeanu, A. Iamnitchi, S. Vazhkudai, “Beyond Music Sharing: An Evaluation of Peer-to-Peer Data Dissemination Techniques in Large Scientific Collaborations,” *Journal of Grid Computing*, Vol. 7, No. 1, pp. 91-114, March 2009.
5. M. Ripeanu, M. P. Singh, and S. S. Vazhkudai, “Virtual Organizations,” *Guest Editors' Introduction to IEEE Internet Computing: Special Issue on Virtual Organizations*, March/April 2008.
6. S. Vazhkudai, X. Ma, “Recovering Transient Data: Automated On-demand Data Reconstruction and Offloading on Supercomputers,” *Operating Systems Review: Special Issue on File and Storage Systems*, Vol. 41, No. 1, pp. 14-18, January 2007.
7. J.W. Cobb, A. Geist, J.A. Kohl, S.D. Miller, P.F. Peterson, G.G. Pike, M.A. Reuter, T. Swain, S.S. Vazhkudai, N.N. Vijayakumar, “The Neutron Science TeraGrid Gateway, a TeraGrid Science Gateway to Support the Spallation Neutron Source,” *Journal of Concurrency and Computation: Practice and Experience*, Vol. 19, pp. 809-826, 2007.
8. S. Vazhkudai, X. Ma, V. Freeh, J. Strickland, N. Tammineedi, T.A. Simon, S.L. Scott, “Constructing Collaborative Desktop Storage Caches for Large Scientific Datasets,” *ACM Transactions on Storage (TOS)*, Volume 2, No. 3, pp. 221-254, August 2006.
9. S. Vazhkudai, “Distributed Downloads of Bulk, Replicated Grid Data,” *Int'l Journal of Grid Computing*, Vol. 2, pp. 31-42, 2004.
10. S. Vazhkudai, J. Schopf, “Using Regression Techniques to Predict Large Data Transfers,” *Int'l Journal of High Performance Computing Applications-Special Issue on Grid Computing: Infrastructure and Applications*, Volume 17, No. 3, pp. 249-268, Fall 2003.
11. S. Vazhkudai, J.M. Syed, P.T. Maginnis, “PODOS - The Design and Implementation of a Performance Oriented Linux Cluster,” *Journal of Future Generation Computer Systems - Special Issue on Cluster Computing*, Volume 18, Issue 3, pp. 335-352, January 2002.

### **Conferences/Workshops**

12. D. Tiwari, S. Boboila, S.S. Vazhkudai, Y. Kim, X. Ma, P.J. Desnoyers and Y. Solihin, “Active Flash: Towards Energy-Efficient, In-Situ Data Analytics on Extreme-Scale Machines”, *Proceedings of the 11th USENIX Conference on File and Storage Technologies (FAST '13)*, San Jose, California, February 2013. (Acceptance Rate: 24/127=18.8%)
13. D. Tiwari, S.S. Vazhkudai, Y. Kim, X. Ma, S. Boboila, P.J. Desnoyers, “Reducing Data Movement Cost using Energy-Efficient Active Computation on SSD”, *Proceedings of the USENIX Workshop on Power-Aware Computing and Systems (HotPower'12, co-located with OSDI'12)*, Hollywood, CA,

October 2012. (10/39=25.6%)

14. A. Khasymski, M.M. Rafique, A.R. Butt, S.S. Vazhkudai, D.S. Nikolopoulos, "On the Use of GPUs in Realizing Cost-Effective Distributed RAID," *Proceedings of the IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'12)*, Washington, D.C., August 2012. (35%)
15. C. Wang, S.S. Vazhkudai, X. Ma, F. Meng, Y. Kim, C. Engelmann, "NVMalloc: Exposing an Aggregate SSD Store as a Memory Partition in Extreme-Scale Machines," *Proceedings of the 26th IEEE Int'l Parallel & Distributed Processing Symposium (IPDPS 2012)*, Shanghai, China, May 2012. (118/568=20.7%)
16. S. Boboila, Y. Kim, S.S. Vazhkudai, P.J. Desnoyers, G. Shipman, "Active Flash: Out-of-core Data Analytics on Flash Storage," *Proceedings of the 28th IEEE Conference on Mass Storage Systems and Technologies (MSST 2012)*, Monterey, CA, April 2012. (14/57=24.5%)
17. S. Boboila, Y. Kim, S.S. Vazhkudai, P.J. Desnoyers, G. Shipman, "Active Flash: Performance-Energy Tradeoffs for Out-of-Core Processing on Non-Volatile Memory Devices," *Poster in the Proceedings of the 3rd Non-Volatile Memories Workshop*, San Diego, CA, March 2012.
18. R. Prabhakar, S. S. Vazhkudai, Y. Kim, A.R. Butt, M. Li, M. Kandemir, "Provisioning a Multi-Tiered Data Staging Area for Extreme-Scale Machines," *Proceedings of the 31st Int'l Conference on Distributed Computing Systems (ICDCS 2011)*, Minneapolis, MN, June 2011. (87/565=15.4%)
19. H. Monti, A.R. Butt, S.S. Vazhkudai, "CATCH: A Cloud-based Adaptive Data Transfer Service for HPC," *Proceedings of the 25th IEEE Int'l Parallel & Distributed Processing Symposium (IPDPS 2011)*, Anchorage, AK, May 2011. (112/512=19.6%)
20. M. Li, S. S. Vazhkudai, A.R. Butt, F. Meng, X. Ma, Y. Kim, C. Engelmann, G. Shipman, "Functional Partitioning to Optimize End-to-End Performance on Many-core Architectures," *Proceedings of Supercomputing 2010 (SC10): 23rd IEEE/ACM Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, New Orleans, LA, November 2010. (51/253=20.1%)
21. Y. Jiao, G. Granroth, C. Griffin, A. Savici, S. Campbell, S.S. Vazhkudai, M. Hagen, S.D. Miller, "Massively Parallel Algorithms for Extreme-Scale Neutron Data Analysis," *Proceedings of the 5th American Conference on Neutron Scattering (ACNS)*, Ottawa, Ontario, Canada, June 2010.
22. S. S. Vazhkudai, M. A. Reuter, J. A. Kohl, S. D. Miller, S. Ren, M. L. Green, "On the Orchestration of the SNS Reduction Job Workflow," *Proceedings of New Opportunities for Better User Group Software (NOBUGS 2010)*, Gatlinburg, TN, October 2010.
23. H. Monti, A.R. Butt, S.S. Vazhkudai, "Reconciling Scratch Space Consumption, Exposure, and Volatility to Achieve Timely Staging of Job Input Data," *Proceedings of the 24th IEEE Int'l Parallel & Distributed Processing Symposium (IPDPS 2010)*, Atlanta, GA, April 2010. (127/527=24.1%)
24. S.D. Miller, K.W. Herwig, S. Ren, S.S. Vazhkudai, P. Jemian, S. Luitz, A.A. Salnikov, I. Gappnenko, T. Proffen, P. Lewis, M.L. Green, "Data Management and Science at DOE BES User Facilities – Past, Present, and Future," *Proceedings of the Scidac 2009*, June 2009.
25. C. Wang, Z. Zhang, X. Ma, S.S. Vazhkudai, F. Mueller, "Improving the Availability of Supercomputer Job Input Data Using Temporal Replication," *Proceedings of the Int'l Supercomputing Conference*, Hamburg, Germany, June 2009.
26. H. Monti, A.R. Butt, S.S. Vazhkudai, "Scratch as a Cache: Rethinking HPC Center Scratch Storage," *Proceedings of the 23rd ACM Int'l Conference on Supercomputing (ICS'09)*, Yorktown Heights, NY, June 2009. (47/191=25%)
27. M.L. Green, S.D. Miller, S.S. Vazhkudai, J.R. Trater, "Doing Your Science While You're in Orbit," *Proceedings of the Int'l Conference on Neutron Scattering*, Knoxville, TN, May 2009.
28. H. Monti, A.R. Butt, S.S. Vazhkudai, "Just-in-time Staging of Large Input Data for Supercomputing Jobs," *Proceedings of Petascale Data Storage Workshop, Supercomputing 2008*, Austin, TX, November 2008.
29. V. E. Lynch, M. L. Chen, J. W. Cobb, E. Farhi, J. A. Kohl, S. D. Miller, P. F. Peterson, M. A. Reuter, J. A. Travieso S. S. Vazhkudai, "SNS-NSTG Collaborative Software Development," *Proceedings of New Opportunities for Better User Group Software (NOBUGS 2008)*, ANSTO, Australia, Nov 3-5, 2008.
30. C. Wang, Z. Zhang, S. S. Vazhkudai, X. Ma, F. Mueller, "On-the-fly Recovery of Job Input Data in Supercomputers," *Proceedings of 37th Int'l Conference on Parallel Processing (ICPP-08)*, Portland, OR, September 2008. (81/263=30.8%)

31. S.A. Kiswany, M. Ripeanu, S. S. Vazhkudai, A. Gharaibeh, "stdchk: A Checkpoint Storage System for Desktop Grid Computing," *Proceedings of the 28<sup>th</sup> Int'l Conference on Distributed Computing Systems (ICDCS 2008)*, Beijing, China, June 2008. (102/638=16%)
32. H. Monti, A.R. Butt, S.S. Vazhkudai, "Timely Offloading of Result-Data in HPC Centers," *Proceedings of the 22<sup>nd</sup> ACM Int'l Conference on Supercomputing (ICS'08)*, Kos, Greece, June 2008. (37/140=26%)
33. H. Monti, A.R. Butt, S.S. Vazhkudai, "A Result-Data offloading Service for HPC Centers," *Proceedings of Petascale Data Storage Workshop, Supercomputing 2007*, Reno, NV, November 2007. (54.6%)
34. Z. Zhang, C. Wang, S. S. Vazhkudai, X. Ma, G. Pike, F. Mueller, J.W. Cobb, "Optimizing Center Performance through Coordinated Data Staging, Scheduling and Recovery," *Proceedings of Supercomputing 2007 (SC07): Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, Reno, NV, November 2007. (54/268=20.1%)
35. S.A. Kiswany, A. Bahramshahry, H. Ghasemi, M. Ripeanu, S.S. Vazhkudai, "A High-Performance GridFTP Server at Desktop Cost," *Poster in Supercomputing 2007 (SC07): Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, Reno, NV, November 2007.
36. S. S. Vazhkudai, J.A. Kohl and J. Schwidder, "A Java-based Science Portal for Neutron Scattering Experiments," *Proceedings of Principles and Practices of Programming in Java (PPPJ 2007)*, Lisbon, Portugal, September, 2007. (19/58=33%)
37. S.A. Kiswany, M. Ripeanu, A. Iamnitchi, S. Vazhkudai, "Are Peer-to-Peer Data Dissemination Techniques Viable in Today's Data Intensive Scientific Collaborations?," *Proceedings of the 13th Int'l European Conference on Parallel and Distributed Computing*, August 2007. (89/333=26.7%)
38. S. Vazhkudai, D. Thain, X. Ma V. Freeh, "Positioning Dynamic Storage Caches for Transient Data," *Proceedings of the Int'l Workshop on High Performance I/O Techniques and Deployment of Very Large Scale I/O Systems (HiperIO'06)*, Barcelona, Spain, September 2006.
39. X. Ma, S. Vazhkudai, V. Freeh, T.A. Simon, T. Yang, S.L. Scott, "Coupling Prefix Caching and Collective Downloads for Remote Data Access," *Proceedings of the Int'l Conference on Supercomputing (ICS06)*, pp. 229-238, Cairns, Australia, June 2006. (37/141=26%)
40. J. Cobb, S. Miller, G. Pike, S. Vazhkudai, M. Hagen, "Neutron Science TeraGrid Gateway," *Proceedings of the TeraGrid'06 Advancing Scientific Discovery*, Indianapolis, IN, June 2006.
41. S. Vazhkudai, X. Ma, V. Freeh, J. Strickland, N. Tammineedi, S.L. Scott, "FreeLoader: Scavenging Desktop Storage Resources for Scientific Data," *Proceedings of Supercomputing 2005 (SC'05): Int'l Conference on High Performance Computing, Networking and Storage*, Seattle, WA, November 2005. (63/260=24.2%)
42. J. Strickland, V. Freeh, X. Ma and S. Vazhkudai, "Governor: Autonomic Throttling for Aggressive Idle Resource Scavenging," *Proceedings of the 2nd IEEE Int'l Conference on Autonomic Computing (ICAC 2005)*, pp. 64-74, Seattle, WA, June 2005. (25/150=16.6%)
43. S. Vazhkudai, "On-demand Grid Storage using Scavenging," *Proceedings of the Session on New Trends in Distributed Data Access*, Las Vegas, Nevada, June 2004.
44. S. Vazhkudai, "Grid Computing—The Hype and The Truth," *Proceedings of the Second Annual Mid-South Computing Conference (Tutorial)*, Little Rock, AR, April 2004.
45. S. Vazhkudai, "Enabling the Co-Allocation of Grid Data Transfers," *Proceedings of the 4th Int'l Workshop on Grid Computing (GRID 2003)*, pp. 44-51, Phoenix, AZ, November 2003. (27/67=40.2%)
46. S. Vazhkudai, J. Schopf, "Using Disk Throughput data in Predictions of End-to-End Grid Transfers," *Proceedings of the 3<sup>rd</sup> Int'l Workshop on Grid Computing (GRID 2002)*, pp. 291-304, Baltimore, MD, November 2002. (28/78=35.8%)
47. S. Vazhkudai, J. Schopf, "Predicting Sporadic Grid Data Transfers," *Proceedings of the 11<sup>th</sup> IEEE Int'l Symposium on High Performance Distributed Computing (HPDC-11)*, pp. 188-196, Edinburgh, Scotland, July 2002. (114/439=26%)
48. S. Vazhkudai, J. Schopf, I. Foster, "Predicting the Performance of Wide-Area Data Transfers," *Proceedings of the 16<sup>th</sup> Int'l Parallel and Distributed Processing Symposium (IPDPS 2002)*, Fort Lauderdale, FL, April 2002. (98/258=38%)
49. S. Vazhkudai, S. Tuecke, I. Foster, "Replica Selection in the Globus Data Grid," *Proceedings of the IEEE Int'l Conference on Cluster Computing and the Grid (CCGRID 2001)*, pp. 106-113, Brisbane, Australia, May 2001. (48/126=38.1%)

50. R. Buyya, S. Vazhkudai, "Compute Power Market: Towards a Market-Oriented Grid," *Proceedings of the IEEE Session on Global Computing on Personal Devices*, pp. 574-581, Brisbane, Australia, May 2001.
51. S. Vazhkudai, G.V. Laszewski, "A Greedy Grid—The Grid Economic Engine Directive," *Proceedings of the IEEE Workshop on Internet Computing and E-Commerce*, San Francisco, CA, April 2001.
52. S. Vazhkudai, S. Tuecke "A Storage Broker for the Globus Environment—A ClassAd Based Implementation," *Poster in Supercomputing 2000*, Dallas, TX, Nov 2000.
53. S. Vazhkudai, P.T. Maginnis, "The PODOS File System—Exploiting the High-Speed Communication Subsystem," *Proceedings of the IEEE Int'l Workshop on Cluster Computing Technologies, Environments, & Applications*, pp. 453-459, Las Vegas, NV, June 2000.
54. S. Vazhkudai, H.C. Cunningham, "A Reusable Software Framework for Distributed Decision-Making Protocols," *Proceedings of the Workshop on Distributed Objects in Computational Science (DOCS'2000)*, pp. 867-873, Las Vegas, NV, June 2000.
55. S. Vazhkudai, P.T. Maginnis, "A High Performance Communication Subsystem for PODOS," *Proceedings of the First IEEE Int'l Conference on Cluster Computing*, pp. 81-91, Melbourne, Australia, December 1999. (60%)

### **Technical Reports**

56. S.A. Kiswany, S. S. Vazhkudai, H. Monti, E. Vairavanathan, M. Ripeanu, "stdchk: A Checkpoint Storage System for the HPC Ecosystem," *ORNL Technical Report*, 2010.
57. S.A. Kiswany, M. Ripeanu, S. S. Vazhkudai, "Aggregate Memory as an Intermediate Checkpoint Storage Device," *ORNL Technical Report 013521*, November 2008.
58. V.W. Freeh, X. Ma, S. Vazhkudai, J. Strickland, "Controlling Impact While Aggressively Scavenging Idle Resources," *ORNL Technical Report*, Oct 2006.
59. V. Freeh, X. Ma, J. Strickland, S. Vazhkudai, "Synergetic Resource Stealing: We Promise It Will Not Hurt Much," *Technical Report P05-123434*, ORNL, May 2005.
60. S. Vazhkudai, "Orchestrating Bulk Data Movement in Grid Environments," *Technical Report R04-121317*, ORNL, September 2004.
61. S. Vazhkudai, "Bulk Data Transfer Forecasts and the Implications to Grid Scheduling," *Ph.D. Dissertation*, University of Mississippi, May 2003.
62. S. Vazhkudai, "Distributed Linux: Evolutionary Steps," *Masters Thesis, Technical Report CISE TR 1998-22*, University of Mississippi, December 1998.

### **Presentations**

1. X. Ma and S.S. Vazhkudai, "Checkpoint I/O for Peta-scale Applications," *Fault Tolerance Workshop*, Dagstuhl, Germany, May 2009.
2. S. S. Vazhkudai, "stdchk: A Checkpoint Storage System for HPC Applications," *TeraGrid Blue Waters Petascale Fault Tolerance Workshop*, March 2009.
3. S. Vazhkudai, "Spallation Neutron Science: State-of-the-art in neutron scattering," *Open Science Grid Meeting*, San Diego, CA, March 2007.
4. S. Vazhkudai, "Optimizing End-User Data Delivery Using Storage Virtualization," *Systems Group Seminar, Department of Computer Science and Engineering, Ohio State University*, October 2006.
5. S. Vazhkudai, "IO Virtualization: Robust Storage Management in the Machine-Room and Beyond," *Virtualization in HPC*, Nashville, TN, September 2006.
6. X. Ma, S. Vazhkudai, V. Freeh, J. Strickland, N. Tammineedi, S.L. Scott, "FreeLoader: Desktop Storage Scavenging," *Intel*, Champaign, IL, December 2005.
7. S. Vazhkudai, S. Studham, "Storage Research at ORNL," *DOE/NSF HEC-IWG Workshop*, Grapevine TX. Aug 16, 2005.
8. S. Vazhkudai, "Towards a Data Management Infrastructure for Spallation Neutron Source," *North Carolina State University*, Raleigh, May 2004.
9. S. Vazhkudai, "Transfer and Storage Management for Grids," *University of Tennessee*, Knoxville, February 2004.
10. S. Vazhkudai, "Peeking into 21<sup>st</sup> Century Distributed Data Intensive Science," at *Vanderbilt University/Univ of Nebraska-Lincoln/Univ of Alabama-Birmingham/Cigital Corporation/Axiom Corporation/Pittsburgh Supercomputing Center*, April-July 2003.
11. S. Vazhkudai, "Server Selection in Data Grids," *University of Southern California*, CA, April 2001.



12. S. Vazhkudai, "A Matchmaking Approach to Replica Selection," *Fermilab*, IL, January 2001.
13. S. Vazhkudai, "The Design and Evolution of Communication in PODOS," *Atlanta Linux Showcase*, GA, October 1999.
14. S. Vazhkudai, "Transmission Group based Communication for PODOS," *Linux World Expo Conference*, CA, August 1999.

#### **Book Chapters**

1. S.S. Vazhkudai, A.R. Butt, X. Ma, "Distributed Storage Systems for Data Intensive Computing," in *Data Intensive Distributed Computing: Challenges and Solutions for Large-scale Information Management*, Editor: Tevfik Kosar, IGI Global Books, January 2012. (ISBN: 9781615209712, DOI: 10.4018/978-1-61520-971-2)

#### **8. PROGRAM DEVELOPMENT ACTIVITIES and GRANTS**

During my nine years at ORNL, I have obtained several external grants from NSF, NIH, and LDRD/SEED awards as highlighted above. The following is a complete listing of my grants.

##### *Current:*

1. S.S. Vazhkudai, X. Ma, D.K. Panda, "Collaborative Research: Dynamic Staging Architecture for Accelerating I/O Pipelines," *NSF High-End Computing University Research Activity (HECURA)*, CCF-0937827, \$672,223, 04/2010-03/2013 (Role: PI).
2. X. Ma, S.S. Vazhkudai, R. Gunasekaran, G. Shipman, "I/O Coordination to Improve Application Performance Stability on Exa-scale Platforms," *DOE ORNL SEED Funds*, \$180,000, FY 2012 (Role: Co-PI).
3. A. Passian, S.S. Vazhkudai, R.H. Farahi, M. Parang, "Physical Sciences-Oncology Center Data Coordination Center (PSOC-DCC)," *NIH-SAIC*, \$956,000, FY 2012-2015 (Role: Co-PI).

##### *Completed:*

1. S.S. Vazhkudai, R. Gillen, D.E. Bernholdt, "Evaluating the Role of Cloud Computing for Scientific Discovery," *DOE ORNL Ultrascale LDRD*, \$655K, FY 2010-2011 (Role: PI).
2. C. Engelmann, S.S. Vazhkudai, "Soft-Error Resilience for Future-Generation High-Performance Computing Systems," *DOE ORNL LDRD*, \$650K, FY 2010-2011 (Role: Co-PI).
3. Y. Jiao, E. Ferragut, S.S. Vazhkudai, M. Hagen, S. Miller, C. Griffin, "Massively Parallel Algorithms for Scalable Exascale Data Analysis," *DOE ORNL LDRD*, \$650K, FY 2010-2011 (Role: Co-PI).
4. X. Ma, Y. Zhou, V.W. Freeh, S. Vazhkudai, "Application-adaptive I/O Stack for High End Computing," *NSF HECURA*, CCF-0621470, \$266K FY 2007-2009 (Role: Co-PI).
5. S. Vazhkudai, X. Ma, J.W. Cobb, G. Pike, "Storage Virtualization: An Integrated Approach to Machine-Room Storage Management," *DOE ORNL LDRD*, \$600K, FY 2007-2008 (Role: PI).
6. C. Engelmann, S.L. Scott, S. Vazhkudai, "Reliability, availability and serviceability (RAS) for terascale computing," *ORNL LDRD*, \$576K, FY 2005-2006 (Role: Co-PI).
7. S. Vazhkudai, G.A. Geist, "A Neutron Science Portal Framework to facilitate Remote Access to SNS Data and Computation," *DOE ORNL LDRD*, \$360K, FY 2004-2005 (Role: PI).

#### **9. AWARDS**

- Significant Event Award (Nov 2012) – Oak Ridge National Laboratory
- Outstanding Mentor Award (Feb 2008) - from the Oak Ridge Institute for Science and Education.
- Doctoral Fellowship (Spring 2003) - U of M (*Awarded to 10 candidates each year*).
- Ph.D. Dissertation Fellowship (2001 - 2002) - Argonne National Laboratory (*Data Grid Research*).
- Wallace Givens Fellowship (Summer 2000) - Argonne National Laboratory (*Awarded to 3 candidates from a group of ~ 40 students*).

#### **10. PROFESSIONAL ACTIVITIES (selected)**

- *Program Committee*: Cluster'12, Networks Architecture and Storage'12, Cluster Computing and Grid'12, Networked Data Management'11, Data Intensive Distributed Computing'11, Resilience'11, Data Aware Distributed Computing'10, Int'l Parallel and Distributed Processing Symposium'09, NAS'09, Principles and Practices of Programming in Java'08, GRID'07, Int'l Conference on Distributed Computing Systems'07, Supercomputing'06.

- *External Reviewer:* Eurosys'12, Journal of Cluster Computing'12, Transactions on Parallel and Distributed Systems'10, Transactions on Storage'10, CCGRID'09, SC'07, TPDS'06, Journal of Grid Computing'05.
- *Proposal Panel Review:* NSF CCF'11, DOE SBIR (2011, 2009), DOE Early Career Awards 2006.
- *Edited* IEEE Internet Computing Journal Special Issue March 2008.

#### **11. PERSONNEL MANAGEMENT and ADVISEES**

Over the course of the past nine years at ORNL, I have had several direct reports in the form of postdoc, post-masters, and several interns every year. I have also managed several research subcontracts to university faculty researchers and as part of this process I have mentored several graduate students.

- *Postdoc:* Chao Wang (UT/ORNL)
- *Post Masters Staff:* Tyler Simon (ORNL)
- *Ph.D. Students Co-advised:* Fei Meng (North Carolina State University, NCSU), Henry Monti (Virginia Tech, VT), Aleksandr Khasymski (VT), Zhe Zhang (NCSU; graduated; now at IBM Research), Samer Kiswany (University of British Columbia), Chao Wang (NCSU; now at ORNL), Ramya Prabhakar (Pennsylvania State University, PSU; graduated; now at NetApp), Simona Boboila (Northeastern University, NEU), Hyogi Sim (VT)
- *Masters Students Co-advised:* Puranjoy Buttarjee (VT; graduated; now at Amazon), Nikhil Komawar (VT), Nandan Tammineedi (NCSU; graduated; now at Yahoo); Jonathan Strickland (NCSU; graduated)
- *Summer Interns:* Devesh Tiwari (NCSU), Simona Boboila (NEU), Ramya Prabhakar (PSU), Sethuraman Subbiah (NCSU), Xin Chen (TN Tech), Baoqiang Yan (Olemiss), Nandan Tammineedi (NCSU), Zhe Zhang (NCSU), Chao Wang (NCSU)